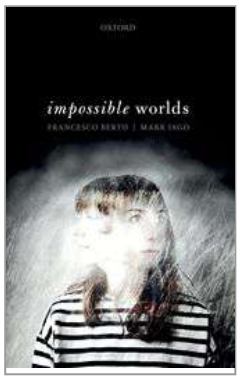

Impossible Worlds
Francesco Berto and Mark Jago

Print publication date: 2019
Print ISBN-13: 9780198812791
Published to Oxford Scholarship Online: August 2019
DOI: 10.1093/oso/9780198812791.001.0001

Ersatz Modal Realism

Francesco Berto
Mark Jago

DOI:10.1093/oso/9780198812791.003.0003

Abstract and Keywords

Ersatz possible worlds can be understood as maximal states of affairs; maximal properties; recombinations of actual bits of reality; as maps; or as entities built from propositions or sentences. The question was: can these approaches be extended to include impossible worlds? The states of affairs approach can, with some modification, accommodate impossible worlds. The property approach too can, with some modification, be extended to impossible worlds. It is argued that the extended approach is best viewed as a form of linguistic ersatzism. The combinatorial faces the question: what *are* recombinations, metaphysically speaking? This approach collapses into one of the others. Map ersatzism does not seem general enough to accommodate all the impossibilities. The most promising approach is linguistic ersatzism. The chapter discusses an issue all ersatz accounts face: the problem of *aliens*.

Keywords: ersatz possible worlds, recombinations, property approach, map ersatzism, linguistic ersatzism, alien properties

3.1 Classifying Ersatz Theories

We are using ‘ersatz world’ in contrast to ‘genuine world’. In §2.2, we understood a genuine world as one which represents the existence of an F by having a real F as a part. So ersatz worlds are those that represent the existence of an F some other way.

That way of drawing the distinction between genuine and ersatz worlds cuts across the distinction between concrete and abstract worlds. Genuine possible

worlds may line up with the concrete worlds, and ersatz possible worlds with the abstract ones. But when we want to talk about impossible worlds as well as possible worlds, these distinctions come apart (§2.2).

Our notion of an ersatz world is essentially a negative one: these are worlds which do not represent in the way genuine worlds do. So how do they represent? Different answers give us different ersatz theories. In §2.4, we introduced a number of ways (in general) to represent *that A*:

(STATE) By using a state of affairs of a certain kind.

(PROPERTY) By using a property reality would have, were things such that *A*;

(COMBINATORIAL) By taking objects and properties which, if recombined in a given way, would make it the case that *A*;

(p.74) (MAP) By using a map, picture, or image which depicts things being such that *A*;

(PROPOSITIONAL) By using the proposition or some other content-carrying entity *that A*;

(LINGUISTIC) By using bits of language, whose meaning is *that A*;

(PRIMITIVE) By taking the relevant representation to be a basic, unanalysable feature.

We've already discussed PRIMITIVE (in §2.7), which we treated as a theory separate from ersatzism. In the rest of this chapter, we'll discuss the remaining options. We'll explore how the various views work in the case of possible worlds, before investigating whether they can also accommodate impossible worlds. In several cases, we'll have to go beyond what the defenders of these views say, since they often don't explicitly discuss impossible worlds. We'll then indicate how we think the views in question may or should be extended to include impossible worlds.

Throughout our discussion, we'll talk of *worlds*, by which we'll normally mean worlds other than the reality around us. As explained in §2.4, ersatz theorists normally include an actualized ersatz world, which represents things being as they in fact are. This move allows them to analyse modal notions purely in terms of ersatz worlds. (Without an ersatz actual world, they'll need to say that *A* is possible iff there's an ersatz-or-genuine world such that *A*, which is a bit clunky. But nothing really hangs on this.)

3.2 Maximal States of Affairs

Plantinga (1970, 1974, 1976) develops a view on which possible worlds are maximal possible states of affairs. A state of affairs is maximal when it either includes or precludes every other state of affairs. Here, both *inclusion* and *preclusion* are understood modally: *s includes t* when, necessarily, if *s* obtains then so does *t*; and *s precludes* **(p.75)** *t* when, necessarily, if *s* obtains then *t* does not. Intuitively, a maximal state of affairs could not have any states of affairs added to it, without becoming impossible.

This approach, then, makes no attempt at a reductive definition of possibility or necessity. To understand 'maximal state of affairs', we first need to understand which states of affairs can obtain with which, and which they cannot obtain without. Modality is taken as a primitive concept.

A key idea is that a state of affairs can exist without *obtaining*. The merely possible worlds are those maximal states of affairs that exist, but do not obtain. Our world – the reality around us – determines the unique obtaining maximal state of affairs. Non-obtaining states of affairs are 'ontologically inert', in the sense that, if the state of affairs *that Fa* does not obtain, then it is not the case that *a* is *F*. We'll discuss just what this might mean at the end of this section.

As a consequence, the approach is *actualist*: all the states of affairs it posits actually exist. When we say 'non-actual world', we mean a state of affairs which exists, but does not obtain. Non-actual worlds represent an *F* by including a state of affairs *that a is F*, for some *a* or other. Such an *a* is represented as being *F* by that state of affairs, and hence by that world. But if that state of affairs does not obtain, *a* is not *F*. That's why we're classifying this view as an ersatz (non-genuine) account of worlds.

Can this approach accommodate impossible worlds? As far as we know, there is little discussion on this in the literature, one exception being Vander Laan 1997. To investigate the issue, we need first to amend the initial approach. Recall that maximality is defined in terms of either including or precluding every other state, with both notions defined modally. Now suppose we add a contradictory state of affairs, *s*, which represents that $A \wedge \neg A$. As *s* cannot possibly obtain, by definition it precludes all states of affairs (including itself), and includes no state of affairs (not even itself). It is maximal, but certainly should not be counted as a possible world.

We do better by thinking of states of affairs structurally, rather than (just) modally. Let us suppose we have some kind of grouping **(p.76)** operation on states of affairs, which makes a plurality into a unity. Think of this as a conjunction operation: if states $s_1, \ldots, s_n$ are grouped as an entity *s*, then *s* is the conjunctive state of affairs whose conjuncts are s_1 to s_n. We might, but need not, treat *s* as a sort of mereological sum or fusion, $s_1 \sqcup \cdots \sqcup s_n$.

This operation would give the notion of inclusion we need: the conjunctive state includes its conjuncts, but no more. Next, take a pair of states s_1, s_2 to be *incompatible* when it is not possible for them to obtain together. A maximal state is one which, for every incompatible pair of possible states s_1, s_2, includes either s_1 or s_2. (Alternatively, we could work with both negative and positive of states of affairs, as Barker and Jago (2012) do, and take a maximal state to be one which includes one state of each pair, *that A, that* $\neg A$.) We then identify the possible worlds with the maximal possible states of affairs, and the impossible worlds with all other states of affairs.

There are thus two ways in which a state of affairs may count as an impossible world. It may be a state that could not possibly obtain (an impossible state), or it may be a non-maximal state. (It may be both.) The former are (metaphysically) inconsistent worlds, the latter (metaphysically) incomplete worlds. We should flag that, on these definitions, an incomplete world may obtain, as part of a larger possible world. The state of affairs *that Charlie's tail is wagging* obtains, as part of the actual world, and hence is possible. But we nevertheless identify that state with an incomplete world, because it could not possibly be all there is to reality.

(We could alter our definitions so that only *non-obtaining* non-maximal states count as incomplete impossible worlds. Then, it would be a contingent matter whether some state *s* counts as a world. That seems strange to us. But since the matter is largely definitional, nothing much hangs on it.)

What does it mean to say a state of affairs exists, but does not obtain? We can't make much sense of the idea. The primary metaphysical role for states of affairs is in accounting for the ontology of predication. States of affairs provide an answer to the question, 'what is it for particular *a* to possess property *F*?' We might **(p.77)** understand the state of affairs *that a is F* as being (in some way) composed of *a* and *F* (Armstrong 1997, 2004). When property and particular are composed in that way, *a* is *F*. That's precisely what it is for *a* to possess *F*. But then, the state of affairs *that a is F* cannot exist without *a* being *F*. So, on Armstrong's approach, a state of affairs obtains just in case it exists.

That argument relies on Armstrong's specific theory of states of affairs. But in general, it is hard to see how some other approach would avoid this problem. If one determines the identity and existence of *that a is F*, but allows that *a* is not *F*, then one undermines the main argument for believing in states of affairs.

There is an independent argument against allowing states of affairs to exist without obtaining (Jago 2018b, §6.2). Those who believe in states of affairs usually take them to be *truthmakers* for the corresponding propositions. The truthmaker for the proposition ⟨*a* is *F*⟩ is the state of affairs *that a is F*. Usually, this is understood as the claim that the existence of the state of affairs *that a is F*

is what makes ⟨*a* is *F*⟩ true. But we cannot say that if states of affairs may exist without obtaining. If *that a is F* exists but does not obtain, then *a* is not *F*, so ⟨*a* is *F*⟩ is not true and hence has no truthmaker.

Instead, the defender of non-obtaining states of affairs will say that a state of affairs makes a proposition true only when it obtains. But this move is problematic. Take a statement of the view in question,

(3.1) There exist states of affairs which do not obtain.

As an existential claim, this is made true (if it is true) by the very things it claims to exist. (Just about everyone agrees with this, even if they do not subscribe to truthmaker theory in general.) The only candidate truthmakers for (3.1) are states of affairs which do not obtain. But this conflicts with the principle above, on which non-obtaining states of affairs do no act as truthmakers. So, we do not think existent but non-obtaining states of affairs are coherent.

Overall, we don't see much hope for overcoming these issues. The notion of non-obtaining states of affairs is absolutely central to the **(p.78)** approach. So if it can't be put on a good metaphysical footing – and we don't see how it can – then the approach is doomed. However, there are other approaches in the vicinity, the propositional and linguistic approaches, which seem to capture the advantages of the states of affairs approach. We'll discuss these in §3.6 below.

3.3 Property Ersatzism

Possible worlds are ways that things could be or have been. Ways are properties. Thus, possible worlds are properties. So says *property ersatzism*. (This view is sometimes called *Stalnakerian realism* or *Stalnakerian ersatzism*, after Stalnaker (1976a); see also Forrest 1986 and Bigelow and Pargetter 1990.) A (non-actual) possible world is identified with the property that would be instantiated, were that world actualized. Such properties must entail a specific pattern of property-instantiations. That is, instantiating a world-property must entail all the specific matters of fact associated with the world in question. As Divers (2002, 177) notes, the approach hasn't been developed greatly, even in the case of possible worlds. So some of what we present in this section is our take on the most promising way to develop the approach.

We might think of a world-property as a big distributional property. Suppose we think of reality in the Humean way, as a pattern of matter, or point-sized particles, scattered over spacetime (Lewis 1986b). Then, a world-property would be a property which specifies a total distribution of matter (or point-sized particles) over spacetime.

This approach to world-properties is rather limited, however. It makes the assumption (usually called *Humean supervenience*: see Lewis 1994) that, given the facts about how matter is distributed locally, all the other facts follow. That's

questionable, even assuming classical physics, and probably wrong given what we know about quantum physics. (Ladyman and Ross (2007) argue at length against the compatibility of Humean supervenience with quantum physics.) **(p. 79)** We mention the idea merely to illustrate what a world-property might look like.

An alternative, and more flexible, approach is to take world-properties to be big conjunctive properties. We can make sense of the idea in terms of an operation on properties, forming a unified property from a plurality. (In §3.1, we mentioned a similar operation on states of affairs. As in that case, the operation on properties may, but need not, be identified with mereological sum or fusion. If we make that identification, then conjunctive properties are fusions of their conjuncts.)

On this approach, the conjoined properties are instantiated by whatever instantiates their conjunction: something possesses *being both F and G* when that very same thing is *F* and also *G*. World-properties are instantiated by reality as a whole, when they are instantiated at all. So, on this view, each conjunct of the conjunctive world-property must be a property possibly possessed by reality as a whole. These will be properties like *being such that Charlie is wagging her tail*. If instantiated, the corresponding state of affairs, *that reality is such that Charlie is wagging her tail*, straightforwardly entails the state of affairs *that Charlie is wagging her tail*.

The details of this approach will be much as they were for the maximal states of affairs account (§3.1). The difference here is that we identify non-actual worlds with uninstantiated properties, rather than non-obtaining states of affairs. On this approach, worlds are actual when they are instantiated (rather than when they obtain). A key advantage of this approach is that it is easier to make metaphysical sense of uninstantiated properties than it is of non-obtaining states of affairs.

When we turn to impossible worlds, the moves we made in §3.1 are attractive options here too. A property *F* is maximal when, for each incompatible pair F_1, F_2 of properties possibly possessed by reality, *F* includes either F_1 or F_2 as a conjunct. (Alternatively, we might work with negative properties, taking F_2 to be the negation of F_1.) We then identify possible worlds with maximal possible properties of reality. Impossible worlds are identified with all other properties of **(p.80)** reality: those that are non-maximal, plus those that reality could not possibly instantiate. (The approach is thus reductive about worlds, but non-reductive about which worlds are possible.)

Understood this way, property ersatzism can be viewed a form of linguistic ersatzism (§3.6). Specific properties of reality, *being such that A*, can be thought of as sentence-like or proposition-like entities, carrying the content *that A*. They

can be conjoined, and perhaps negated too. So we will defer further discussion of this approach until §3.6, where we will discuss the issues faced by linguistic ersatzism.

3.4 Combinatorial Ersatzism

Combinatorial ersatzism comes in various shapes, but all of its variants agree that possible worlds should be understood in terms of recombinations of bits of actual reality. These can be actual individuals and actually instantiated properties, or obtaining simple states of affairs. Examples of the strategy are found in Quine 1969 and Cresswell 1972. We will focus on Armstrong's (1989, 1997) exemplary version, which he develops in terms of recombinations of actual particulars and universals (properties and relations). Key to this idea is Armstrong's notion of *sparse universals*: fundamental properties like *charge, spin*, and *mass*, as opposed to derivative ones like *being a penguin* and *being an* X Factor *contestant*. We can't freely recombine any property with any particular. Electrons can't possibly be red, and you can't possibly have spin ½. But, says Armstrong, we can freely recombine the sparse properties with the simple individuals.

Armstrong's aim is to give a reductive account of modality. He spells out which (actual) particulars and universals may be freely recombined, and spells out what it is to recombine them. All of this is done without bringing in modal notions. Possible worlds are then identified with recombinations of the selected particulars and universals, and modality is analysed in the usual way in terms of possible worlds.

(p.81) On this picture, each possible world consists of a rearrangement of fundamental bits of reality. But possible worlds represent more than that: they also represent facts about penguins, *X-Factor* contestants, and the like. On Armstrong's view, all the non-fundamental facts represented by a world come for free, given facts about what's fundamental (according to that world). What a possible world *is*, according to Armstrong, is a rearrangement of fundamental stuff. But what it represents is whatever follows, metaphysically, from those facts, rearranged thus-and-so.

There are a number of issues with the approach, which we'll mention only briefly here. First, it's clear that not every fundamental bit of reality exists necessarily. This or that quark (or whatever) need not have existed. So merely rearranging actual, fundamental stuff won't give the right results. Instead, we have to think in terms of rearrangements which potentially leave some bits out.

By the same token, it seems that there could have been (fundamental) things which don't in fact exist. There could have been more quarks (or whatever) than in fact there are. No rearrangement of actual stuff produces those non-actual entities. So we have to think in terms of rearrangements which potentially include extra stuff (as well as potentially leaving some bits out). Those extra bits

are *alien entities*: entities which are possible, but which don't actually exist. They're a real problem for this approach (very likely, for any actualist approach). After all, what are we supposed to be arranging? If it's actual stuff, plus some extra bits, then we seem to have quantified over non-actual entities, and the actualist endeavour is over. Moreover, if we are allowed to quantify over non-actual entities, why not instead go for genuine modal realism (§2.2)? This is a serious problem, but it's a problem any ersatz account must face. We'll set it aside for now, and return to it in §3.7.

The issue we want to investigate here is: just what *are* recombinations, metaphysically speaking? We'll argue that the recombination approach is destined to collapse into one of the other approaches we've discussed. Suppose rearrangements are (or consist of) states of affairs. Do the non-actual states of affairs obtain or not? If we say **(p.82)** they don't obtain, then we have the states of affairs account from §3.1. If they do obtain (perhaps on the grounds that, for states of affairs, *obtaining* amounts to *existing*), then we have a form of genuine realism. (Or we may have no theory at all. If possible worlds consist of genuine states of affairs, then modal space as a whole will contain contradictory states of affairs. It's hard to see how this approach can avoid outright triviality, wherein every sentence is treated as being true *simpliciter* (Jago 2018b).)

Rearrangements might be treated not as states of affairs, but rather as ersatz replacements for them. These might be propositions, or some other kind of property-and-particular-containing representational entity. That approach will be a form of propositional or linguistic ersatzism (to be discussed in §3.6). Armstrong's view is that the rearrangements do not in fact exist; but it is nevertheless convenient and acceptable to speak as if they do (Armstrong 1989, 49–51). As Sider (2005) argues, this is a form of fictionalism (§2.8). Possible-worlds talk should be understood as talk about a fictional ontology of recombinations. So, however we understand the ontology of recombinations, we don't have a genuinely distinct account of what possible worlds are.

It's also worth flagging a specific difficulty for Armstrong's fictional recombinations. If the fiction talks of all the recombinations at once, then it's at serious risk of inconsistency. In that fiction, there will exist pairs of logically incompatible states of affairs. Their existence renders that fiction logically inconsistent, and this in turn renders the entire account inconsistent (Jago 2016). The problem is avoided if there's a separate fiction for each recombination. But then, there's no way to analyse iterated modality, as when we say, it's possibly necessary that such-and-such.

Combinatorialism isn't best viewed as a theory of what possible worlds are. Rather, it's a theory of the extent of the space of possible worlds. How do things change when we include impossible worlds? We are not aware of any proposals in the literature, and so what follows are our own suggestions. It seems clear

that, if all the rearrangements correspond to possible worlds, then we **(p.83)** can't make sense of impossible worlds in terms of those same rearrangements. One option is to make space in the theory for rearrangements which aren't licensed by Armstrong's restriction to sparse universals. We allow the derivative properties to 'float free' of their fundamental grounds. So, rather than rearranging facts about tables by manipulating fundamental facts about (say) the arrangement of matter in spacetime, we could think in terms of the property *being a table* itself. We may rearrange properties to arrive at a (representation of a) massless table, for example.

This approach will generate (representations of) impossibilities. We can recover the possible worlds as those that conform to acceptable recombinations of Armstrong's sparse universals. To do that, we'll need enough *bridge principles*, from fundamental to derivative facts. These 'vertical' principles tell us how the derivative facts are *grounded* in the fundamental facts. They should tell us, for example, that there's no fundamental recombination to support our massless table, and hence that that recombination gives us an impossible world. In general, when one of these grounding principles is broken, we have an impossible world.

This notion of *grounding* – a form of metaphysical dependence, of the kind we briefly discussed in §1.3 – is a much-debated concept in contemporary metaphysics (Fine 2001, 2012b,c, Schaffer 2009). If sense can be made of the notion using non-modal vocabulary, then the approach suggested here may offer a reductive and actualist-friendly analysis of modality.

3.5 Map Ersatzism

Map erstzism takes worlds to represent in much the same way as a realistic painting, a map, or architect's scale model does (Lewis 1986b, chapter 3). Consider how a map works. It has a limited vocabulary of symbols, arranged spatially, in a way that mimics the spatial arrangement of the bits of reality thereby represented. Maps typically use *above* and *below* (on the map) to represent *north of* **(p.84)** and *south of* (in reality), so that a tree-symbol above a lake-symbol represents a situation in which there are trees to the north of a lake. The distance between the symbols on the map, together with the map's scale, represents the distance in reality between trees and lake.

Map-based representations have some very useful representational properties. Suppose, in our map with a tree-symbol above the lake-symbol, there's also a church-symbol under the lake-symbol. Then the map represents the trees as being north of the church. And it represents the church as being south of both the trees and the lake. It does this for free. If we'd represented the tree-lake and lake-church relations descriptively, using sentences, then we'd have to do some work to infer that the trees are north of the church. And similarly, if we'd represented those facts linguistically using 'north', then we'd have to do some

further work to infer that the church is south of the trees. Map-based representations build this inferential work automatically into the way they represent.

This feature makes map-based representations great candidates where what is represented is closed under (some amount of) logical consequence. Block (1983), for example, argues that cognitive data suggests that (at least some) mental representations are image-like, rather than linguistic. (See also Kosslyn and Pomerantz 1977, Paivio 1986, and Pinker 1980.)

But by the same token, this feature makes it hard for map-based representations to represent certain *im*possibilities. Just take the impossible situation where the trees are north of the lake, but the lake isn't south of the trees. How could a map represent that? Escher was ingenious in finding ways to depict many impossible situations pictorially, but the applicability of these or any other pictorial techniques is fairly limited. How would a map or picture depict an explicitly contradictory situation, say the one described by Graham Priest in *Sylvan's Box* (Priest 1997b), in which the narrator discovers a box that is both empty and not empty? (We discuss *Sylvan's Box* in detail in §11.3.) How would we even begin to depict the impossible situation in which Fermat's Last Theorem turns out false?

(p.85) There are options here, to be sure. A map may contain a special *not like this* symbol, with maps featuring that symbol interpreted as the global negation of what they would otherwise represent. We may further introduce a conjunction operation on maps, with the resulting map representing the conjunction of what each map individually represents. Conjunction plus negation will allow us to represent explicit impossibilities, for sure. But we get there using linguistic, not pictorial, techniques. It seems likely that any form of map or pictorial ersatzism will need to resort to non-pictorial, linguistic ways of representing, if it aims to represent all the desired possibilities and impossibilities. But then why not go for linguistic representation across the board, and do away with any worries about how the pictorial and linguistic elements are supposed to interact?

3.6 Propositional and Linguistic Ersatzism

Propositional ersatzism identifies worlds with sets of propositions. Linguistic ersatzism is similar, except it builds worlds from sentences of some chosen language, the *worldmaking language*, rather than propositions. Linguistic forms of ersatzism go back to Carnap's (1947) theory of *state descriptions*, and both Adams (1974) and Jeffrey (1965) talk of *world-stories*, understood as sets of propositions. (They don't have to be sets. They could be some other kind of construction from propositions or sentences. All that matters is that we can take all the propositions or sentences making for a world together, treating them as a single entity. We'll work with sets.)

One nice feature of such worlds is the definition of truth, relative to a world. It is true that A, relative to world w, just in case the proposition *that* A, or the worldmaking sentence which means *that* A, is a member of world w. In this kind of definition, we have a phrase 'that A' in both the definiens and the definiendum. But it isn't circular. We're defining the truth of an English sentence, relative to a world, either in terms of a proposition or of a sentence of some other **(p.86)** language, the worldmaking language. In the latter case, this amounts to a translation from the worldmaking language into English.

Just what properties worlds have depends on our theory of propositions or our construction of the worldmaking language. Let's consider propositions first. One prominent theory of propositions identifies them with sets of possible worlds (Lewis 1986b, Stalnaker 1984). But that approach is ruled out, if we first identify worlds with sets of propositions. Indeed, that is one motivation for rejecting propositional ersatzism altogether. Modal semantics often works with sets of worlds, and it's very natural to identify these sets with propositions, as we saw in Chapter 1.

Whether propositions really are sets of worlds is moot (Jago 2015, 2018b, Merricks 2015). An alternative is the *Russellian* account, after Russell (1904/1980), which identifies (atomic) propositions with ordered sequences (or tuples) of particulars, properties, and relations. King (1995, 1996, 2011), Salmon (1986, 2005), and Soames (1987, 2008) all defend a variant of this view. The proposition *that Charlie is wagging her tail* is the sequence 《*wagging*, Charlie, Charlie's tail》. (We'll use the notation '《 》' for sequences and the more familiar '⟨ ⟩' for propositions.)

Logically complex propositions are treated as sequences which include (the semantic values of) logical connectives. We might identify these values with truth-functions, familiar from truth-table semantics (although we don't need to make this identification). Let NEG, CONJ, and DISJ be the truth-functions corresponding to *negation, conjunction*, and *disjunction*, respectively. Then we can treat negated, conjunctive, and disjunctive propositions as follows:

$$\begin{aligned} \langle \neg A \rangle &= \text{《 } \mathrm{NEG}, \langle A \rangle \text{ 》} \\ \langle A \wedge B \rangle &= \text{《 } \mathrm{CONJ}, \langle A \rangle, \langle B \rangle \text{ 》} \\ \langle A \vee B \rangle &= \text{《 } \mathrm{DISJ}, \langle A \rangle, \langle B \rangle \text{ 》} \end{aligned}$$

Since order matters to sequences, the structure of these propositions is (relatively) fine-grained, and allows for distinct but logically equivalent propositions. Even very closely related propositions, such **(p.87)** as $\langle A \wedge B \rangle$ and $\langle B \wedge A \rangle$, will be treated as distinct entities, on this approach. We'll say more about the Russellian approach below.

Yet another approach is to take propositions to be metaphysically primitive, unanalysable entities (Merricks 2015). We discussed the corresponding primitivist view about worlds in §2.7. Many of the same considerations raised there apply also to primitivism about propositions, and so we won't discuss the view any further.

Let's turn to linguistic ersatzism. Linguistic ersatzism is similar to propositional ersatzism, but with sentences in place of propositions. Worlds can be taken as sets of sentences. A theory along these lines must first specify the language to which those sentences belong. This is what is called the worldmaking language. The choice of language will be important. To see why, suppose we adopted English as our worldmaking language. Then, a world *w* represents that *A* iff it contains the English sentence '*A*'. And so, on this analysis, it is possible that *A* iff there's a possible world which contains the English sentence '*A*'. To determine whether this is the case, we will need to ask, which sets of English sentences are *compossible* (that is, possibly true all together)? But answering this question will depend, in part, on whether '*A*' represents a possible situation. So we seem to have a circular analysis of possibility.

To avoid circularity, the worldmaking language and the theorist's language (English, in this case) must differ. One option is to use a *Lagadonian* language (Lewis 1986b, 145–6, following Carnap 1947), in which each particular is a name for itself and each property and relation is a predicate denoting itself. Infinitary logical connectives allow for sentences of infinite length (Divers 2002, 180).

The approach achieves little (goes one objection) if it requires that we first analyse which sets of sentences count as possible worlds. We do not have a reductive account of modality, of the kind genuine realists claim to have (§2.2). This is one of two objections Lewis (1986b) raises against linguistic ersatzism. (We'll discuss his second objection in §3.7.) The approach applies equally to propositional ersatzism, if the propositions in question are fine-grained as on the Russellian account.

(p.88) In response, a defender of propositional or linguistic ersatzism may say that she never intended her account to be fully reductive. Its benefits lie elsewhere, she may claim. Sider (2002), for instance, argues that

> Unless modal consistency can be reduced in some way, linguistic-ersatz worlds cannot be used in a reductive analysis of modality, on pain of circularity. But the linguistic ersatzist can accept this limitation. The reduction of worlds to language still has a point, for it allows us to reduce all talk of worlds—which runs far beyond that which can be said utilizing merely the modal operators—to talk of possibility and necessity. As for these, they may one day be reduced in some way that does not involve worlds, or they may remain primitive. (Sider 2002, 282)

> A reduction of talk of possibilia which employs primitive possibility and necessity is nevertheless valuable since talk of possibilia runs beyond what can be said in the language of quantified modal logic.
>
> (Sider 2002, 306)

That's a concessive response. A more ambitious propositional or linguistic ersatzer might attempt to offer a fully reductive account of modality, by saying, in non-modal terms, which sets of propositions or sentences count as possible worlds. The difficulty here is that non-modal notions appear to be either too narrow, or too broad, to pick out metaphysical modality. To see why, consider those sets of propositions or sentences that are both logically consistent and closed under (classical) logical consequence. These are *maximal consistent* sets, and correspond to (classical) logically possible worlds. But not all of these represent (genuine, metaphysical, etc.) possibilities. There will be worlds where this apple is both red all over and green all over, or where there are married bachelors; these things aren't logically inconsistent. Lewis (1986b) argues, on these grounds, that this purely syntactic way of delimiting possible worlds gets the modal facts wrong, by delivering too broad a conception of what's possible.

The ambitious, reductive approach is difficult. Nevertheless, here's an attempt, which builds on the approach sketched at the end of §3.4. (We'll focus now on linguistic ersatzism, but a propositional **(p.89)** ersatzer could say similar things.) The metaphysically possible worlds are those logically possible worlds (the maximal consistent sets of sentences) which respect *grounding relationships*. Let's unpack this. *Grounding* is a metaphysical dependency relationship (§1.3) which holds between facts or states of affairs (and perhaps between other kinds of entity), where the grounded state of affairs obtains *in virtue of* the ground. The fact that it's now 20°C in Nottingham grounds the (more general) fact that it's now hot in Nottingham. If physicalism about the mind is true, then the fact that Anna's brain is in such-and-such state grounds the fact that she's feeling hot. And that fact in turn grounds the fact that the proposition ⟨Anna is feeling hot⟩ is true.

A set of sentences S respects grounding relationships as follows. Suppose S contains sentences $A_1, \ldots, A_n$, which represent states of affairs $s_1, \ldots, s_n$. Suppose also that $s_1, \ldots, s_n$ together ground a further state of affairs s. Then any sentence A which represents s must be a member of S, too. The idea is then that, when such sets of worldmaking sentences are consistent and closed under logical consequence, they (and only they) count as possible worlds.

This approach will work only if *grounding* is itself a non-modal notion. As we saw in §1.3, Fine (2001, 2012b) argues persuasively that grounding is irreducible to modal notions. Grounding has modal consequences – if state s_1 grounds s_2, then it is necessary that s_2 obtains if s_1 obtains – but it is not itself defined in modal

terms. Indeed, it might be that grounding is a primitive, undefinable metaphysical notion. One of us attempts a reductive, non-modal analysis of grounding in Jago 2018a. Whether this approach will ultimately work out is beyond the scope of our discussion here. But if it does, then a fully reductive account of *possible world* to non-modal notions is on the cards for linguistic (and propositional) ersatzism.

3.7 Alien Entities

All versions of ersatz realism face a worry with *alien entities*: things that do not (actually) exist, but could have. Franz could have had an **(p.90)** older sister (but, in fact, doesn't). How is his possible older sister to be represented by an ersatz possible world? One way is by naming her. Let's call her 'Franzista'. But how do we give that name meaning? Not by associating it directly with an individual (as we do at a naming ceremony, for example), since, by hypothesis, there actually is no such person. We could instead give 'Franzista' its meaning by associating it with the description, 'Franz's older sister'. But this gives us strange results. Franz might not have existed. So we should expect the following situation to be possible: Franzista exists, but Franz does not. So, in that possible situation, Franzista is not Franz's sister. But this situation is contradictory if 'Franzista' means *Franz's sister*.

A better approach is to forget about naming Franz's merely possible older sister, and instead *describe* our target possible world. We can characterize that world as follows: *there are distinct individuals* x_1, x_2, …, *each distinct from all actual individuals, such that* x_1 *is such-and-such,* x_2 *is such-and-such*, …, x_i *is Franz's older sister*, …. Here, we describe a possible situation in which there exists some older sister of Franz's. We can allow that that person may have had no siblings herself in some other possible situation.

A problem remains. How are we to capture this further possible situation, in which Franz's (merely possible) older sister has no siblings? Call the possible world just described, containing Franz and a Franz-sister, w_1. Let w_2 be a Franz-less world, in which Franz's sister from w_1 has no siblings. What guarantees that we are talking about the same person, Franz's possible sister, in both worlds? For all we've said, nothing does. So we have no way to say, thinking about Franz's sister in w_1, that *she* could have had no siblings.

One response to the problem is to use *counterpart theory* (Lewis 1968, 1971). This is a feature of Lewisian genuine modal realism, though we did not get into it in §2.2. Counterparts capture *de re* ways an entity could have been. If x and y are counterparts, then x could have been like y, and vice versa. In our ersatz setting, we can describe some x_i in w_1, and describe some x_j in w_2, and then say that x_1 in w_1 and x_j in w_2 are counterparts. One problem here is that the counterpart relation is, well, a relation, and relations relate entities. **(p.91)** Unless one is a Meinongian, if it's true that x_i and x_j are related, then x_i and x_j

exist. (Note that the theory doesn't say that x_i and x_j are counterparts *according to world* w_1 or any other world. It's a claim of the theory from the viewpoint of all modal space.) But, by hypothesis, neither x_i nor x_j exist. So this can't be the right approach.

The problem of aliens gets even tougher to solve when we turn our attention to merely possible properties (Lewis 1986b, chapter 3). Presumably, there are properties that could have been instantiated but are not. Perhaps there are fundamental uninstantiated properties. They pose a serious problem for any ersatz approach. At least we have the means to describe Franz's merely possible older sister. We do so using the predicates 'older than' and 'sister of', which get their meaning through being associated with the corresponding relations. But in the case of alien fundamental properties, we don't even have the vocabulary to describe them.

Sider (2002) shows how to solve these problems. On his approach, rather than describing each possible world one by one, we describe all of them in one go. (For simplicity, we'll consider the case of alien individuals only.) We say: *there are distinct individuals* x_1, x_2, ..., *each distinct from all actual individuals, and worlds* w_1, w_2, ..., *such that, in* w_1, ..., *and in* w_2, ..., *and* .... Here, we quantify over merely possible individuals all in one go, and then describe how each is in each of the worlds which represent that individual. This is the *Ersatz pluriverse sentence*. It is false, of course, since (by actualist lights) there are no merely possible individuals. Its function is to represent a space of possibilities, just as (standard) linguistic ersatzism does. The ersatz pluriverse sentence makes clear when distinct worlds w_i and w_j are talking about the same possible particular. That is the great advantage of Sider's approach. Jago (2013c) raises a worry for Sider's approach, which we won't discuss here. This said, we conclude by saying that we take a form of linguistic ersatzism endowed with grounding relations to be the most promising ersatz account. One advantage of linguistic ersatzism is that it helps with an insidious objection to impossible worlds: the *compositionality objection*, which we discuss in §8.5.

(p.92) Chapter Summary

We can understand ersatz possible worlds as maximal states of affairs; maximal properties; recombinations of actual bits of reality; as maps; or as entities built from propositions or sentences (§3.1). Our question was: can these approaches be extended to include impossible worlds? The states of affairs approach can, with some modification, accommodate impossible worlds; yet we found the concept of a non-obtaining state of affairs hard to sustain (§3.2). The property approach too can, with some modification, be extended to impossible worlds. We then argued (§3.3) that the extended approach is best viewed as a form of linguistic ersatzism.

The combinatorial faces the question: what *are* recombinations, metaphysically speaking? We argued that, however the question is answered, this approach collapses into one of the others (§3.4). Map ersatzism does not seem general enough to accommodate all the impossibilities our theories require of impossible worlds (§3.5). The most promising approach seems to be propositional or linguistic ersatzism, of which, we prefer the linguistic variant (§3.6). Finally, we discussed an issue all ersatz accounts face: the problem of *alien entities* (§3.7).